On the Voting Time of the Deterministic Majority Process

Dominik Kaaser
University of Salzburg
dominik@cosy.sbg.ac.at

Frederik Mallmann-Trenn
École Normale Supérieure
Simon Fraser University
fmallman@sfu.ca

Emanuele Natale
Sapienza Università di Roma
natale@di.uniroma1.it

Abstract

In the deterministic binary majority process we are given a simple graph where each node has one out of two initial opinions. In every round, every node adopts the majority opinion among its neighbors. By using a potential argument first discovered by Goles and Olivos (1980), it is known that this process always converges in $\mathrm{O}(|E|)$ rounds to a two-periodic state in which every node either keeps its opinion or changes it in every round.

It has been shown by Frischknecht, Keller, and Wattenhofer (2013) that the $\mathrm{O}(|E|)$ bound on the convergence time of the deterministic binary majority process is indeed tight even for dense graphs. However, in many graphs such as the complete graph, from any initial opinion assignment, the process converges in just a constant number of rounds.

By carefully exploiting the structure of the potential function by Goles and Olivos (1980), we derive a new upper bound on the voting time of the deterministic binary majority process that accounts for such exceptional cases. We show that it is possible to identify certain modules of a graph G in order to obtain a new graph G^{Δ} with the property that the worst-case convergence time of G^{Δ} is an upper bound on that of G. Moreover, even though our upper bound can be computed in linear time, we show that, given an integer k, it is NP-hard to decide whether there exists an initial opinion assignment for which it takes more than k rounds to converge to the two-periodic state.

1. Introduction

We study the *deterministic binary majority process* which is defined as follows. We are given a graph $G = (V, E)$ where each node has one out of two opinions. The process runs in discrete rounds where each node in every round computes and adopts the majority opinion among all of its neighbors.

It is known that this deterministic process always converges to a two-periodic state. The *convergence time* of a given graph for a given initial opinion assignment is the time it takes until the two-periodic state is reached. In this work we give bounds on the

voting time, which is the maximum convergence time over all possible initial opinion assignments.

The deterministic binary majority process has widespread applications in the study of *influence networks* in distributed computing [FKW13], distributed databases [Gif79], sensor networks [BTV09], the competition of opinions in social networks [MT14b], social behavior in game theory [DP94], chemical reaction networks [AAE07], neural and automata networks [GM90], and cells' behavior in biology [CCN12]. Variants of the deterministic binary majority process have been used in the area of distributed community detection [RAK07, KPS13, CG10], where the voting time is essentially the convergence time of the proposed community-detection protocols.

Among its many probabilistic variants that have been previously considered, plenty of work concerns *randomized voting* where in each step every node is allowed to contact a random sample of its neighbors and updates its current opinion according to the majority opinion in that sample [AF02, BMPS05, CEOR13, DW83, HL75, HP01, Lig85, LN07, MT14a, Oli12].

In an algorithmic game theoretic setting, the deterministic binary majority process can be seen as the simplest *discrete preference games* [CKO13]. In this game theoretic perspective, the existence of so-called monopolies has been investigated [ACF+14]. A *monopoly* in a graph is a set of nodes which start with the same opinion, causing all other nodes to eventually adopt this opinion. In the distributed computing area, a lot of research has been done to find small monopolies, see for example [Pel02]. It has also been shown that there exist families of graphs with constant-size monopolies [Ber01]. More recently, classes of graphs which do not have small monopolies have been investigated [Pel14].

Many of these results relate to the voting time of the deterministic binary majority process. It was proven independently by Goles and Olivos [GO80], and Poljak and Sůra [PS83] with the same potential function argument that the deterministic binary majority process always converges to a two-periodic state. They later (independently) refined and generalized the potential function argument in several directions [GCFSP85, PT86, GO88, Gol89]. Their proof was popularized in the *Puzzled* columns of Communications of the ACM [Win08a, Win08b]. Recently, the same problem has been studied on infinite graphs w.r.t. a given probability distribution on the initial opinion assignments [BCO+14]. In [TT13], the authors provide a bound on the number of times a node in a given bounded-degree graph changes its opinion. Both [BCO+14] and [TT13] also investigate the probability that in the two-periodic state all nodes hold the same opinion.

As for the maximum time it takes for the process to converge over all initial opinion assignments, Frischknecht et al. [FKW13] note that the potential argument by Goles et al. [GO80, PS83, Win08b] can be used to prove an $\mathrm{O}(|E|)$ upper bound. They furthermore show that this upper bound is tight in general, by designing a class of graphs in which the deterministic binary majority process takes at least $\Omega\left(|V|^2\right)$ rounds to converge from a given initial opinion assignment. This construction has later been extended to prove lower bounds for weighted and multi-edges graphs by Keller et al. [KPW14].

A lot of attention has been given to the two-periodic state to which the deterministic

binary majority process converges to. However, besides the $\mathrm{O}(|E|)$ upper bound that follows from the result by Goles et al. [GO80, PS83, Win08b], no further upper bound on the voting time that holds for any initial opinion assignment has been proved. Still, one can observe that in many graphs the voting time is much smaller than $\mathrm{O}(|E|)$, e.g., the voting time of the complete graph is one.

We show that for the deterministic binary majority process the question whether the voting time is greater than a given number is NP-hard. While for many generalizations of the deterministic binary majority process many decision problems are known to be NP-hard, at the best of our knowledge this is the first NP-hardness proof that does not require any additional mechanisms besides the bare majority rule of the deterministic binary majority process. However, as we show in the rest of the paper, it is possible to obtain upper bounds on the voting time which can be computed in linear time. A module of a graph is a subset of vertices S such that for each pair of nodes $u, v \in S$ it holds that $N(u) \setminus S = N(v) \setminus S$. By carefully exploiting the structure of the potential function by Goles et al. and leveraging the particular behavior that certain modules of the graph exhibit in the deterministic binary majority process, we are able to prove that the voting time of a graph can be bounded by that of a smaller graph that can be constructed in linear time by contracting suitable vertices.

We obtain a new upper bound that asymptotically improves on the previous $\mathrm{O}(|E|)$ bound on graph classes which are characterized by a high number of modules that are either cliques or independent sets. For instance, the Turán graph $T(n, r)$ is the graph formed by partitioning a set of n vertices into r subsets, with sizes as equal as possible, and connecting two vertices by an edge whenever they belong to different subsets. For the convergence time of the Turán graph $T(n, r)$ we obtain an $\mathrm{O}(r^2)$ bound, compared to the previously best known bound of $\mathrm{O}(n^2)$. Also, for the convergence time of full d-ary trees we get an $\mathrm{O}(|V|/d)$ bound, compared to $\mathrm{O}(|V|)$ originating from the $\mathrm{O}(|E|)$ bounds. Our bound relies on a well-known graph contraction technique, e.g., see the notion of *identical vertices* in [SSKc13].

1.1. Preliminaries

We are given a connected graph $G = (V, E)$ and an initial opinion assignment which we define now.

Definition 1. *An opinion assignment f_t in round $t \geq 0$ is a function $f_t : V \to \{0, 1\}$ which assigns for each $v \in V$ an opinion with*

$$f_t(v) = \begin{cases} 1 & \textit{if } v \textit{ has opinion } 1 \textit{ at time } t \\ 0 & \textit{if } v \textit{ has opinion } 0 \textit{ at time } t. \end{cases}$$

We will also denote opinion 1 *as* white *and opinion* 0 *as* black. *The opinion assignment at time $t = 0$ is called* initial opinion assignment.

The deterministic binary majority process can be defined as follows. Let v be an arbitrary but fixed vertex and $N(v)$ the set of neighbors of v. To compute $f_{t+1}(v)$ the

node v computes the majority opinion of all of its neighbors in $N(v)$. In the case of a tie the node behaves *lazily*, that is, v stays with its own opinion. Otherwise, there is a *clear majority* and the node adopts the majority opinion. This leads to the following definition.

Definition 2. *Let $G = (V, E)$ be a graph and let f_0 be an initial opinion assignment such that $f_0 : V \to \{0, 1\}$. The deterministic binary majority process is the series of opinion assignments that satisfies the following rule.*

$$f_{t+1}(v) = \begin{cases} 0 & \text{if } |\{u \in N(v) : f_t(u) = 0\}| > |\{u \in N(v) : f_t(u) = 1\}| \\ 1 & \text{if } |\{u \in N(v) : f_t(u) = 0\}| < |\{u \in N(v) : f_t(u) = 1\}| \\ f_t(v) & \text{otherwise} \end{cases}$$

Note that the pair (G, f_0) completely determines the behavior of the system according to the majority process. We next define the main object of this work, namely the voting time.

Definition 3. *Given a graph $G = (V, E)$ and any initial opinion assignment f_0 on V, the* convergence time $\mathfrak{T}$ *of the majority process on G w.r.t. f_0 is $\mathfrak{T} = \mathfrak{T}(G, f_0) = \min \{t : \forall v \; f_{t+2}(v) = f_t(v)\}$. The* voting time *of G is defined as $\max_{f_0 \in \{0,1\}^V} \mathfrak{T}(G, f_0)$.*

Observe that $\mathfrak{T}$ is indeed the number of steps until the process converges to a periodic state. This holds since the process is completely determined by the current opinion assignment. Thus $f_{t+2}(v) = f_t(v)$ also implies that $f_{t+3}(v) = f_{t+1}(v)$ for all nodes v.

In the following we assume without loss of generality that G is connected. For disconnected graphs the deterministic binary majority process runs independently in each connected component. Therefore, the resulting upper bounds on the voting time time can be replaced by the maximum over the corresponding bounds in the individual connected components of G.

1.2. Our Contribution

First we define the *voting time decision problem* VTDP and show that it is NP-complete.

Definition 4 (voting time decision problem, VTDP)**.** *For a given graph G and an integer k, is there an assignment of initial opinions such that the voting time of G is at least k?*

Theorem 1.1. *Given a general simple graph G,* VTDP *is NP-complete.*

Then, in Section 3 we extend known approaches to derive upper bounds on the voting time, which are tight for general graphs. In Section 3.2, we identify the following subsets of nodes that play a crucial role in determining the voting time of the deterministic binary majority process.

Definition 5. *A set of nodes S is called a family if and only if for all pairs of nodes $u, v \in S$ we have $N(u) \setminus \{v\} = N(v) \setminus \{u\}$. We say that a family S is* proper *if $|S| > 1$.*

The set of families of a graph forms a partition of the nodes into equivalence classes. Our main contribution is a proof that the voting time of the deterministic binary majority process is bounded by that of a new graph obtained by contracting its families into one or two nodes, as stated in the following theorem.

Definition 6. *Given a graph $G = (V, E)$, its asymmetric graph $G^\Delta = (V^\Delta, E^\Delta)$ is the subgraph of G induced by the subset $V^\Delta \subseteq V$ constructed by replacing every family of odd-degree non-adjacent nodes with one node and replacing any other proper family with two nodes.*

Theorem 1.2. *Given any initial opinion assignment on a graph $G = (V, E)$, the voting time of the deterministic binary majority process is at most*

$$1 + \min\left\{|E^\Delta| - \frac{|V^\Delta_{odd}|}{2}, \frac{|E^\Delta|}{2} + \frac{|V^\Delta_{even}|}{4} + \frac{7}{4}\cdot|V^\Delta|\right\} \ .$$

Furthermore, this bound can be computed in $\mathrm{O}\left(|E|\right)$ *time.*

As mentioned before, this bound becomes $\mathrm{O}\left(r^2\right)$ for the Turán graph $T(n, r)$ and $\mathrm{O}\left(|V|/d\right)$ for d-ary trees. Finally, in Appendix C of the appendix, we give some insight into the computational properties of the voting time.

2. NP-Completeness

If it was possible to efficiently compute the worst-case voting time, there would have been not much interest in investigating good upper bounds for it. In this section, we show that this is unlikely to be the case. We prove Theorem 1.1 by reducing 3SAT to the voting time decision problem. Given $\Phi \in$ 3SAT, we construct a graph $G = G(\Phi)$ such that the deterministic binary majority process on G simulates the evaluation of Φ. The graph G consists of h layers. The first layer represents an assignment of the variables in Φ, the remaining layers represent Φ and ensure that the assignment of variables in Φ is valid. We will show that if Φ is satisfiable, then there exists an initial assignment of opinions for which the convergence time is exactly $h+1$. If, however, Φ is not satisfiable, then any assignment of opinions will result in a convergence time strictly less than $h+1$. We now give the formal proof.

Reduction

Let $\Phi \in$ 3SAT be a Boolean formula in 3-conjunctive normal form. Let n be the number of variables of Φ. Let m be the number of clauses of Φ. The Boolean formula is of the form

$$\Phi = (l_{1,1} \vee l_{1,2} \vee l_{1,3}) \wedge \cdots \wedge (l_{m,1} \vee l_{m,2} \vee l_{m,3})$$

where $l_{i,j} \in \{x_1, \overline{x}_1, x_2, \overline{x}_2, \cdots, x_n, \overline{x}_n\}$ is a literal for $1 \leq i \leq m$ and $1 \leq j \leq 3$.

We construct a graph G to simulate the evaluation of Φ as follows. Let $\ell = 10 \cdot (m+n) + 1$. The graph consists of several layers. On the first layer, we place so-called

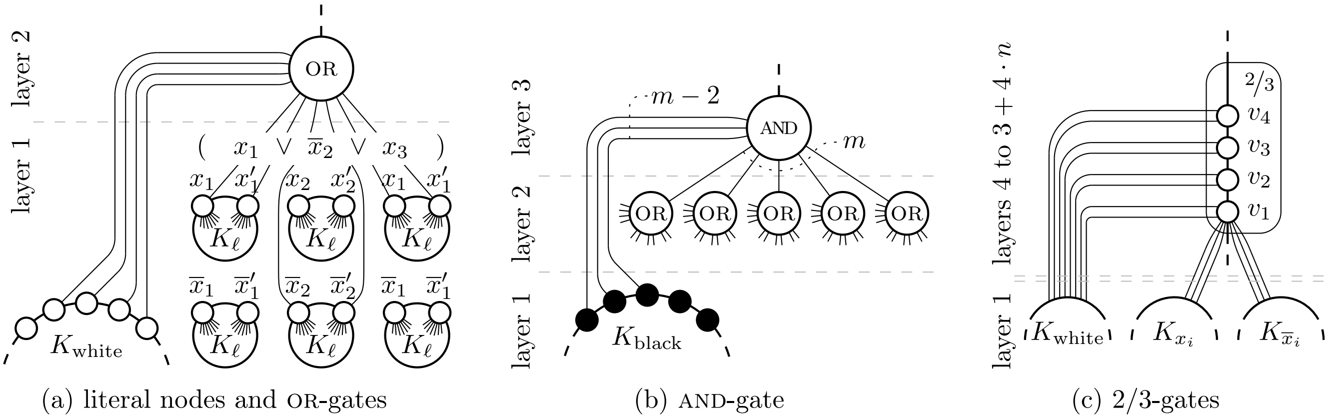

Figure 1: The gates and layers used in the reduction from 3SAT to VTDP.

literal cliques of size ℓ, and on the layers above we place the gates. In our reduction, we use OR-gates, an AND-gate, and 2/3-gates. Each gate consists of one or several nodes. Additionally, we have two so-called *mega-cliques* K_{white} and K_{black} of size ℓ.

Let g be an arbitrary but fixed gate. In the following, we will denote a node on a layer below g that is connected to g as *input node* to g. Additionally, we will denote a node that belongs to g and is connected to another gate on a layer above g as *output node*.

In the following, we assume that opinion 1, white, corresponds to Boolean TRUE and 0, black, corresponds to FALSE. The main idea of the construction is to show that an *activation signal* is transmitted from the bottom up through all layers. If the current assignment of opinions on the literal cliques corresponds to a satisfying assignment of Boolean values to Φ, then the process requires $4 + 4 \cdot n$ steps. The main purpose of the OR-gates and the AND-gate is to evaluate Φ. The 2/3-gates check whether the opinion assignment to literal nodes is valid. That is, we need to enforce that the corresponding literal nodes for x_i and $\overline{x}_i$ are of opposite colors for every variable x_i of Φ. If either this condition is violated and variables x_i exist for which $x_i = \overline{x}_i$ or the current assignment of opinions on the literal cliques does not corresponds to a satisfying assignment of Boolean values to Φ, the construction enforces that the process stops prematurely after strictly fewer than $h + 1$ steps.

We start by giving a detailed description of the gates and the layers used in our construction.

Layer 1 – Literal Cliques. We represent each variable x_i with two cliques, one for x_i and one for $\overline{x}_i$. Each clique has a size of ℓ which is defined above. Note that ℓ is odd. Additionally, we distinguish three so-called *representative nodes* in each of these cliques. Furthermore, we add two cliques of size ℓ to the graph which we call mega-cliques. Intuitively, these mega-cliques represent the Boolean values TRUE and FALSE. We will show that they cannot have the same color in order to achieve a long convergence time. The mega-cliques are used in all other gates.

Layer 2 – Parallel OR-Gates. The OR-gates are placed on layer 2 and consist of one node v which is also the output node. There is one OR-gate for every clause. Fix a clause $(l_{j,1} \vee l_{j,2} \vee l_{j,3})$. Input nodes are three pairs of nodes $(v_1,\ v_1')$, (v_2, v_2'), and (v_3, v_3'), where $(v_1,\ v_1')$ are two representative nodes of the literal clique for $l_{j,1}$, (v_2, v_2') are representatives of $l_{j,2}$, and (v_3, v_3') are representatives of $l_{j,3}$. That is, for each literal in the clause we connect the OR-gate on layer 2 to two of the three representative nodes of the corresponding literal clique on layer 1. The output node v is additionally connected to 4 nodes of the K_{white} mega-clique. Intuitively, we use the OR-gates to verify that for each clause at least one literal is TRUE. All clauses are evaluated simultaneously using an OR-gate for each clause. The OR-gate is shown in Figure 1a.

Layer 3 – AND-Gate. There is exactly one AND-gate on layer 3. This AND-gate consists of one output node denoted u_0, which has the following input nodes. It is connected to every output node of the OR-gates on layer 2 and to $m-2$ distinct nodes of the K_{black} mega-clique. Intuitively, the AND-gate is used to verify that every clause is satisfied.

Layers 4 to $3 + 4 \cdot n$ – 2/3-Gates. The 2/3-gates consist of a path v_1, v_2, v_3, and v_4. Each node of this path is connected to two distinct nodes of the K_{white}. The output node of the gate is v_4. The node v_1 of the first 2/3-gate on layer 4 is connected to the AND-gate on layer 3. The node v_1 of each of the following 2/3-gates is connected to the node v_4 of the previous 2/3-gate. Additionally, the input node of the i-th 2/3-gate is connected to three distinct nodes of the literal clique representing x_i and to three distinct nodes of the literal clique representing $\overline{x}_i$ on layer 1. The output node of the final 2/3-gate is connected to K_{black}. This is shown in Figure 1c. These gates are used to verify that we do not have variables x_i in Φ for which the literal cliques of x_i and $\overline{x}_i$ have the same color. Observe that 2/3-gates span over 4 layers and we have n such 2/3-gates.

Literal cliques, OR-gates, and the AND-gate use only one layer, while 2/3-gates span over 4 layers. Therefore, the total number of layers is $h = 3 + 4 \cdot n$, which results from one layer for the literal cliques, one layer for the OR-gates, one layer for the AND-gate, and $4 \cdot n$ layers containing n concatenated 2/3-gates. Based on above description of G we prove the following lemmas, which are then used to show Theorem 1.1.

Lemma 2.1. *If Φ is satisfiable, then there exists an assignment of opinions such that the convergence time in G is at least $h + 1$.*

Proof. Let $A = (a_1, \ldots, a_n)$ be an assignment of Boolean values to the n variables in Φ which satisfies Φ. We need to show that there exists an opinion assignment on G for which the deterministic binary majority process requires at least $h+1$ steps to converge. In the following, we construct such an opinion assignment.

Let f_A be an initial opinion assignment in the graph G that represents A by initializing the nodes in the literal cliques on layer 1 according to the assignment A as follows. For each literal x_i or $\overline{x}_i$, a_i assigns either TRUE or FALSE to the literal. We denote a literal x_i or $\overline{x}_i$ which is assigned TRUE as *positive* and literals which are assigned FALSE as *negative.*

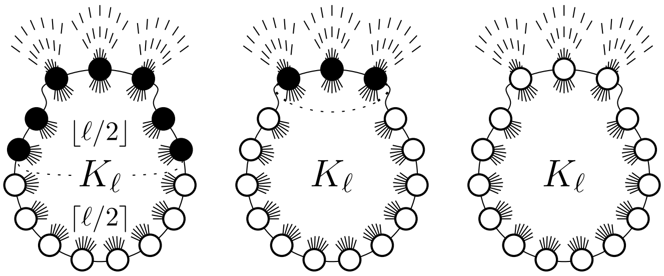

(a) $t = 0$ (b) $t = 1$ (c) $t = 2$

Figure 2: The behavior of the cliques on layer 1. The three top nodes are the representatives.

For the positive literal cliques, we assign the color black to $\lfloor L/2 \rfloor$ nodes including the representative nodes of the clique. The remaining $\lceil L/2 \rceil$ nodes, which do not have any other connections except within the literal clique, are colored white. Negative literal cliques are colored entirely black. Furthermore, we initialize all nodes of the K_{white} and the K_{black} with white and black, respectively. All other nodes, the paths v_1 to v_4 in the 2/3-gates, the output nodes of the OR-gates, and the output node of the AND-gate, are colored black.

The process now behaves as follows.

1. In the first step, all black nodes in every positive literal clique except the representative nodes turn white, since they have $\lceil L/2 \rceil$ white neighbors and only $\lfloor L/2 \rfloor - 1$ black neighbors.

2. The representative nodes of the literal cliques will turn white in the following step. This behavior of the cliques on layer 1 is shown in Figure 2.

3. Additionally to the neighbors in K_{white}, all OR-gates on layer 2 will have at least two white input nodes from representing at least one positive literal clique, since A satisfies Φ. Therefore, the OR-gates will turn white in step 3.

4. Once all OR-gates become white, the AND-gate has a total of m white input nodes that form a clear majority against the $m - 2$ edges to black nodes in K_{black} and the 1 edge to the black node of the first 2/3-gate. Therefore, the AND-gate turns white in step 4.

5. In the following $4 \cdot n$ steps, node after node and gate after gate the 2/3-gates turn white. Once all nodes of the 2/3 gates have turned white, the process stops.

Summing up over all of the above steps, the convergence time w.r.t. the initial opinion assignment f_A is exactly $\mathfrak{T}(G(\Phi), f_A) = 4 + 4 \cdot n = h + 1$. Therefore, the voting time in $G(\Phi)$ for a satisfiable Φ is at least $h + 1$, which yields the lemma. □

It remains to show that if Φ is not satisfiable, then the voting time in G is strictly less than $h + 1$. Recall that the voting time is the maximum of the convergence time over all possible initial opinion assignments.

Lemma 2.2. *If Φ is not satisfiable, then there is no assignment of opinions such that the convergence time in G is at least $h + 1$.*

Before we prove this lemma, we establish several auxiliary lemmas which require the following definitions. Let u_0 denote the output node of the AND-gate. Consider the graph G' induced by the nodes of the AND-gate and the nodes of the 2/3-gates. Let u_i be the node at distance i to u_0 in G'. That is, G' is a path that consists of the $4 \cdot n + 1$ top layers of the graph G and u_i is the i-th node on this path.

Definition 7 (Stable Time)**.** *We define the* stable time $s(v)$ *for any node $v \in V$ to be the first time step such that v does not change its opinion in any subsequent time step $t' > s(v)$ over all possible initial configurations. That is,*

$$s(v) = \min \left\{ t : \forall f_0 \in \{0,1\}^V \;\; \forall t' \geq t \quad f_{t'}(v) = f_t(v) \right\} \quad .$$

Accordingly, let for any subset $V' \subseteq V$ be $s(V')$ defined as $s(V') = \max \{s(v) : v \in V'\}$.

In the following, let V_K be the set of nodes of all cliques in $G(\Phi)$, that is, the nodes contained in the literal cliques and in the mega-cliques on layer 1. Furthermore, let $V_{K^{\mathrm{R}}}$ be the set of representatives of the cliques and $V_{K^-} = V_K \setminus V_{K^{\mathrm{R}}}$. That is, every clique K on layer 1 consists of $K^- \cup K^{\mathrm{R}}$. Finally, let V_{OR} be the set of all output nodes of OR-gates. The following lemma shows that the layers become stable one after the other.

Lemma 2.3. *It takes at most 3 time steps for the layers 1 and 2 consisting of literal cliques and OR-gates to become stable. Precisely, we have 1) $s(V_{K^-}) = 1$, 2) $s(V_{K^{\mathrm{R}}}) = 2$, and 3) $s(V_{\mathrm{OR}}) = 3$.*

Proof. The lower bounds for all three claims follow from the initial opinion assignment f_A described in the proof of Lemma 2.1. We now show the upper bounds. In the following, let f_0 be an arbitrary but fixed initial opinion assignment.

(i) Let K be an arbitrary but fixed clique and let $c \in \{0, 1\}$ be the majority color among the nodes of K. Let furthermore K^- be the set of clique nodes that do not have connections to any other node except within the clique, that is, K^- contains all clique nodes except representatives. Note that all nodes in K^- only have connections to all other nodes in K. Since K is odd and c is the majority color in K, each node $v \in K^-$ with $f_0(v) = c$ will have at least $\lfloor \ell/2 \rfloor$ neighbors out of a total of $\ell - 1$ neighbors colored c. Therefore, each node $v \in K^-$ with $f_0(v) = c$ will keep its color c such that $f_1(v) = c$. However, all other nodes $v' \in K^-$ with $f_0(v') \neq c$ will change their opinion to c, since they have at least $\lceil \ell/2 \rceil$ neighbors out of a total of $\ell - 1$ neighbors colored c, such that $f_1(v') = c$.

By construction and by the size of the clique, ℓ, all nodes $v \in K^-$ have more neighbors in K^- than in $V \setminus K^-$. Therefore, for all consecutive steps $t' \geq 1$, we have for any $v \in K^-$ that $f_{t'}(v) = c$. This holds for all cliques on layer 1, and thus $s(V_{K^-}) \leq 1$.

(ii) Let K be an arbitrary but fixed clique and let $v \in K^{\mathrm{R}}$ be a representative node of K. By construction, v has a majority of its neighbors in K^{-} and hence from (i) we derive $s(v) \leq 2$. Therefore, $s(V_{K^{\mathrm{R}}}) \leq 2$. We also observe that all nodes in K^{R} have the same color after the second step, since the nodes in K^{-} become monochromatic in the first step and these nodes dominate the behavior of the nodes in K^{R}.

(iii) Let v be the output node of an arbitrary but fixed OR-gate in V_{OR}. We observe that all neighbors of v except for one neighbor (the node of the AND-gate u_0) are stable for any time step $t' \geq 2$. By (ii), at time 2 all representatives of any literal x_i have the same color and K_{black} is stable. Therefore, at time 2 an even number of neighbors of v are black and an even number is white. Since the total number of neighbors of v is 10, we observe that u_0 cannot influence $f_{t'}(v)$ for $t' \geq 2$. Moreover, by (i) and (ii) we have at time $t' \geq 2$ that the majority of neighbors having color c does not change and therefore v becomes stable at time 3. Thus $s(V_{\mathrm{OR}}) \leq 3$ holds. □

The above lemma gives bounds on the stable time of layers 1 and 2. In the following, we argue that whenever a node changes its opinion in any step t after time step 3, it will not change its color in any subsequent time step $t' \geq t$ any more. We therefore define the so-called *activation time* of a node $v \in G'$ as follows.

Definition 8 (Activation Time). *Let c be the color of the K_{black} mega-clique at time 2 and let f_0 be an arbitrary but fixed initial opinion assignment. We define the* activation time *of a node $v \in G'$ to be the first time step after time step 3 in which the node v adopts opinion c. That is, $a(v) = \min\{t \geq 3 : f_t(v) = c\}$. If v does not change its color after time step 3 we write $a(v) = 3$.*

We now use above definition to state the following lemma, which describes that every node $u_i \in G'$ with $i \geq 1$ changes its color at most once after time step 3. Note that this covers the nodes of the 2/3-gates.

Lemma 2.4. *Let f_0 be an arbitrary but fixed initial opinion assignment. Let t be the activation time w.r.t. f_0 of the node $u_i \in G'$ with $i \geq 1$ such that $t = a(u_i)$. Then for all $t' \geq t$ we have $f_{t'}(u_i) = f_t(u_i)$.*

Proof. By Lemma 2.3, all nodes $u \in V_{K^{\mathrm{R}}}$ are stable at $t' \geq 2$. We now distinguish two cases.

Case 1: $i \mod 4 \neq 1$. Observe that u_i can only change its color at time $t = a(u_i)$, if it had a different color than K_{white} in the previous round. This holds, since every node u_i with $i \mod 4 \neq 1$ has the same number of connections to K_{white} than to nodes in $V \setminus K_{\text{white}}$. Since furthermore the process behaves lazy, any node u_i which has the same color as K_{white} cannot change its opinion back to the opposite color any more.

Case 2: $i \mod 4 = 1$. The node u_i is a v_1 node of the j-th 2/3-gate with $j = \lceil i/4 \rceil$. Therefore it is connected to three representatives of each literal clique for x_j and $\overline{x}_j$. The literal representatives of x_j and $\overline{x}_j$ are stable at time $t' \geq 2$. Now if x_j and $\overline{x}_j$ have the same color c, then u_i has $6 > |N(u_i)|/2$ edges to nodes of color c. Therefore, the node does not change its color any more after time step 3. That is, we have $a(u_i) = 3$ and also $f_{t'}(u_i) = c$ for any consecutive time step $t' \geq 3$. If, however, x_j and $\overline{x}_j$ do not have the same color, these edge *cancel* each other out and the color of node u_i is determined by u_{i-1}, u_{i+1}, and K_{white}. Therefore, the same argument as in the first case holds. □

In the following we examine the behavior of layer 3 which contains only of the AND-gate. Recall that u_0 is the output node of the AND-gate. The next lemma describes the following fact. The AND-gate u_0 can only change its color in a time step $t \geq 4$ if u_1 changed its color in time step $t-1$. After this change at time t, the node u_0 cannot change its color again.

Lemma 2.5. *Let f_0 be an arbitrary but fixed initial opinion assignment and let furthermore t be the round after node u_1 has been activated such that $t = a(u_1) + 1$. For all consecutive rounds $t' \geq t$ we have $f_{t'}(u_0) = f_t(u_0)$. That is, the* AND*-gate does not change its opinion any more once the node u_1 has become stable.* (see Appendix A)

The following lemma implies that in order to reach a convergence time of $h+1$ the gates on the path $u_0, \ldots, u_k$ in G' have to activate one after the other starting with u_0 at time 4. Recall that $k = 4 \cdot n$.

Lemma 2.6. *Let f_0 be an arbitrary but fixed initial opinion assignment and let $u_i \in G'$ be a node with $0 \leq i \leq k$. If $a(u_i) < i + 4$ w.r.t. f_0, then $\mathfrak{T}(G(\Phi), f_0) < h + 1$.* (see Appendix A)

In the following two lemmas, we enforce that initial opinion assignments which do not represent valid assignments of Boolean values to literal cliques result in premature termination of the deterministic binary majority process in $G(\Phi)$. An assignment is called *illegal* if there exist literal cliques such that the majority of x_i and the majority of $\overline{x}_i$ have the same initial color.

Lemma 2.7. *Let f_I be an illegal initial opinion assignment. The convergence time $\mathfrak{T}(G(\Phi), f_I)$ is strictly less than $h + 1$.* (see Appendix A)

Lemma 2.8. *If after two time steps K_{white} and K_{black} have the same color, the process stops after strictly fewer steps than $h + 1$.*

Proof. Let c be the color of both mega-cliques after two time steps. Note that from Lemma 2.3 we conclude that both cliques are stable at time 2. Therefore u_k activates at most at time 3, that is, $a(u_k) = 3$. By induction, one can show that u_i will activate at most at time $3 + k - i$. Hence u_1 becomes activated at most at time $t = 3 + k - 1 < h$ and u_0 at most at time $t = 3 + k$ which is strictly less than $h + 1$. Since by Lemma 2.3 all other nodes are also stable at time $3 + k < h + 1$ the claim follows. □

We now combine above lemmas and prove Lemma 2.2.

Proof of Lemma 2.2. In the following we assume that K_{white} and K_{black} have opposite colors after the second step, since otherwise the convergence time is less than $h+1$ as shown in Lemma 2.8. W.l.o.g., assume K_{white} is colored white and K_{black} is colored black. Furthermore, we assume that the assignment is legal, since otherwise the convergence time is less than $h+1$ as shown in Lemma 2.7. Finally, we also assume that $u_1, \ldots, u_k$ are initially black, since otherwise the convergence time is less than $h+1$ as shown in Lemma 2.6. Note that this especially covers the node u_1 which we assume to be black at time 4, since otherwise again the convergence time is less than $h+1$ according to Lemma 2.6.

According to the assumption of Lemma 2.2, Φ is not satisfiable. That is, for every possible assignment of Boolean values to the variables in Φ, there exists a clause $(l_1 \vee l_2 \vee l_3)$ where all literals l_1, l_2, and l_3 are FALSE. Therefore, for any legal initial opinion assignment f_0 in $G(\Phi)$, the representative nodes of the corresponding literal cliques will be black at time 2. Consequently, the OR-gate corresponding to that clause will be stable with color black at time 3.

This implies that the AND-gate is black as long as u_4 is black since at least $(m-2)+1+1 > |N(u_0)|/2$ neighbors are black. Since the AND-gate is black, we can only have $a(u_1) = 5$ if $a(u_2) = 4$. According to Lemma 2.6, this results in a convergence time strictly less than $h+1$. Note that if $f_3(u_1) = 1$, then u_2 will be activated at time 4 and again by Lemma 2.6 this yields that the convergence time is less than $h+1$. □

Now we combine Lemma 2.1 and Lemma 2.2 to show Theorem 1.1.

Proof of Theorem 1.1. It is easy to see that VTDP is in NP. Furthermore, we can polynomially reduce 3SAT to VTDP. The correctness proof of the reduction follows from Lemma 2.1 and Lemma 2.2. Therefore, VTDP is NP-complete. □

3. Bounds on the Voting Time

Since the problem is NP hard, we cannot hope to calculate the voting time of a graph efficiently. Nevertheless, in this section we show, that it is possible to obtain non-trivial upper bounds on the voting time that are easy to compute. This section is dedicated to proving our upper bound on the voting time.Theorem 1.2. The main contribution of this theorem is the influence of symmetry which is studied in Section 3.2.

We start by giving a formal version of the potential function argument [GO80, PS83] as conceived in [Win08b]. In the following we assume that each edge in $\{x, y\} \in E$ can be replaced by two directed edges (x, y) and (y, x). The main idea is based on so-called *bad arrows* defined as follows.

Definition 9. *Let $G = (V, E)$ be a graph with initial opinion assignment f_0. Let v denote an arbitrary but fixed node and $u \in N(v)$ a neighbor of v. Let t denote an arbitrary but fixed round. The directed edge (v, u) is called* bad arrow *if and only if the opinion of u in round $t+1$ differs from the opinion of v in round t.*

Intuitively, each of these directed edges (v, u) can be seen as *advice* given from v to u in the voting process. In the case of a bad arrow the advice was not followed by u since it has a different opinion in the following round than v. Observe that each bad arrow is incident at exactly two nodes and thus we say it is *outgoing* in the node at its tail and incoming in the node at its head. An example of such a bad arrow can be seen in Figure 3.

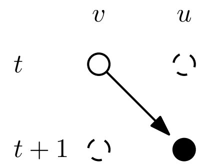

Figure 3: A *bad arrow* from node v to node u in round t.

Theorem 3.1. *Let $G = (V, E)$ be a graph which contains only vertices of odd degree. The voting time of the deterministic binary majority process on G is at most $1 + W$ where W is an upper bound on the initial number of bad arrows for any initial opinion assignment on G. In particular, the voting time of G is at most $2 \cdot |E| + 1$.*

For a proof, see Appendix B. Note that in Theorem 3.1 it is assumed that all nodes of the graph have odd-degree. In the following we show how to remove this assumption. Let in the following V_{even} be the set of even-degree vertices in V and, analogously, V_{odd} be the set of odd-degree vertices. Clearly, $V = V_{\text{even}} \cup V_{\text{odd}}$.

Definition 10. *Let $G = (V, E)$ be a graph. The graph $G^* = (V, E^*)$ is the graph obtained by adding a self loop to every node of even degree in G. More formally,*

$$E^* = E \ \cup \bigcup_{v \in V_{\text{even}}} (v, v) \quad .$$

From the definition it follows that $|E^*| = |E| + |V_{\text{even}}|$.

Theorem 3.2. *The voting time of the deterministic binary majority process on any graph $G = (V, E)$ is at most $1 + W_{bad}$, where W_{bad} is the number of initial bad arrows in G^*.*

For a proof, see Appendix B. The upper bound on the voting time considered in [KPW14] follows from the $2 \cdot |E|$ upper bound on the number of bad arrows of Theorem 3.1. In the following we show that this result can be improved further by a factor of 2 by simply applying the following lemma.

Lemma 3.3. *Let G be a graph. The number of initial bad arrows in G^* can be at most $|E| - |V_{odd}|/2$.*

Therefore, combining Theorem 3.2 and Lemma 3.3 we obtain the following corollary.

Corollary 3.4. *The voting time of the deterministic binary majority process on any graph $G = (V, E)$ is at most $1 + |E| - |V_{odd}|/2$.*

Note that Corollary 3.4 is tight for general graphs up to an additive constant of 1. Indeed, consider a path with an initial opinion assignment on which the opinions alternate except for the last two nodes, which share the same opinion.

3.1. Improved Bounds for Dense Graphs

We observe that Corollary 3.4 is (almost) tight, and it gives us a voting time linear in the number of vertices for sparse graphs where $|E| = \mathrm{O}(|V|)$. However, for dense graphs with $|E| = \Omega\left(|V|^2\right)$ there is room for improvement. Now the main goal in this following subsection is to reduce the dominant term of the voting time even further, which leads us to the following theorem.

Theorem 3.5. *Let $G = (V, E)$ denote a graph. For any initial opinion assignment f_0 on G, the voting time of the deterministic binary majority process is at most $1 + \frac{|E|}{2} + \frac{|V_{\text{even}}|}{4} + \frac{7}{4} \cdot |V|$.* (see Appendix B)

3.2. The Influence of Symmetry

We observe that the deterministic binary majority process is much faster on graphs that exhibit certain types of symmetry, such as the star graph, the complete graph and many other graphs in which several nodes share a common neighborhood. We investigate this feature of the process to further improve the bounds obtained so far. We recall that a set of nodes S is called a *family* if and only if for all nodes $u, v \in S$ we have $N(u) \setminus \{v\} = N(v) \setminus \{u\}$. The key fact is that these nodes of any family will behave in a similar way after the first step.

Definition 11. *Let $\operatorname{fam}(u)$ denote the family u belongs to. We write $u \sim v$ if $\operatorname{fam}(u) = \operatorname{fam}(v)$.*

Lemma 3.6. *The relation $\sim$ defines an equivalence class. In particular, all nodes in the same family either form a clique or are all pairwise non-adjacent, and they all have the same degree in G.* (see Appendix B)

Corollary 3.7. *For any graph G, its asymmetric graph G^{Δ} is well-defined.*

Proof. Thanks to Lemma 3.6, the set of families is a partition of the nodes of G. By construction of G^{Δ}, every family S in G is replaced by one or two nodes in G^{Δ}. Therefore, there is a bijection between the families in G and the corresponding node or pair of nodes in G^{Δ}. Hence G^{Δ} is well-defined. $\square$

We now prove Theorem 1.2.

Proof of Theorem 1.2. Let v and v' be two nodes of the same family $\operatorname{fam}(v) = \operatorname{fam}(v')$, having the same color at time t. Since v and v' observe the same opinions in their respective neighborhood, v and v' will also have the same color anytime after t. It follows that if at some time t there is a bad arrow going from v to some neighbor u (or from u to v), then there will also be a bad arrow from v' to u (or from u to v'). In particular, this implies that whenever the number of bad arrows adjacent to v is decreased by some amount c, also the identical number of bad arrows adjacent to v' will be decrease by the same amount c.

Now recall the proofs of Corollary 3.4 and Theorem 3.5. An estimate of the voting time is obtained by upper bounding the number of bad arrows that can possibly disappear during the process. The main argument is the following. It suffices to only consider the bad-arrows adjacent to v in G^{Δ}, since the corresponding bad arrows adjacent to v' will disappear whenever those adjacent to v do.

Now let v and v' be two nodes with $\mathrm{fam}(v) = \mathrm{fam}(v')$ having a different color at time t. We can divide every such family that contains nodes of different opinions into two sets S_0 and S_1 according to their initial opinion in the first round. Note that all nodes in either set behave identically. In particular, an adjacent bad arrow from a node u to all nodes of either set disappears at the same time. Since there is bijection between the families of G and the pairs of nodes and singletons of G^{Δ}, and by applying Corollary 3.4 and Theorem 3.5 we can bound the voting time by bounding the bad arrows in G^{Δ}. This yields the first part of the claim. Using [CH94], one can obtain the modular composition of G in $\mathrm{O}(|E|)$ time steps. In another $\mathrm{O}(|E|)$ time steps one can select from the modular composition those modules that form a family, using that all nodes of a family have the same degree. Hence, G^{Δ} can be constructed in linear time. □

Acknowledgments. We would like to thank our supervisors Petra Berenbrink, Andrea Clementi, and Robert Elsässer for helpful discussions and important hints.

References

[AAE07] D. Angluin, J. Aspnes, and D. Eisenstat. A Simple Population Protocol for Fast Robust Approximate Majority. In *Proc. DISC '07*, pages 20–32, 2007.

[ACF$^+$14] V. Auletta, I. Caragiannis, D. Ferraioli, C. Galdi, and G. Persiano. Discrete preference games: social influence through coordination, and beyond. *CoRR*, abs/1402.4050, 2014.

[AF02] D. Aldous and J. Fill. Reversible Markov Chains and Random Walks on Graphs, 2002. Unpublished. http://www.stat.berkeley.edu/~aldous/RWG/book.html.

[BCO$^+$14] I. Benjamini, S.-O. Chan, R. O'Donnell, O. Tamuz, and L.-Y. Tan. Convergence, unanimity and disagreement in majority dynamics on unimodular graphs and random graphs. *CoRR*, abs/1405.2486, 2014.

[Ber01] E. Berger. Dynamic Monopolies of Constant Size. *Journal of Combinatorial Theory, Series B*, 83(2):191–200, 2001.

[BMPS05] S. Brahma, S. Macharla, S.P. Pal, and S.K. Singh. Fair Leader Election by Randomized Voting. In *Proc. ICDCIT '05*, pages 22–31, 2005.

[BTV09] F. Bénézit, P. Thiran, and M. Vetterli. Interval consensus: From quantized gossip to voting. In *Proc. ICASSP '09*, pages 3661–3664, 2009.

[CCN12] L. Cardelli and A. Csikász-Nagy. The Cell Cycle Switch Computes Approximate Majority. *Scientific Reports*, 2(656), 2012.

[CEOR13] C. Cooper, R. Elsässer, H. Ono, and T. Radzik. Coalescing Random Walks and Voting on Connected Graphs. *SIAM Journal on Discrete Mathematics*, 27(4):1748–1758, 2013.

[CG10] G. Cordasco and L. Gargano. Community Detection via Semi–Synchronous Label Propagation Algorithms. In *Proc. BASNA '10*, pages 1–8, 2010.

[CH94] A. Cournier and M. Habib. A New Linear Algorithm for Modular Decomposition. In *Proc. 19th Colloquium on Trees in Algebra and Programming (CAAP '94)*, pages 68–84, 1994.

[CKO13] F. Chierichetti, J.M. Kleinberg, and S. Oren. On Discrete Preferences and Coordination. In *Proc. EC '13*, pages 233–250, 2013.

[DP94] X. Deng and C. Papadimitriou. On the Complexity of Cooperative Solution Concepts. *Mathematics of Operations Research*, 19(2):257–266, 1994.

[DW83] P. Donnelly and D. Welsh. Finite particle systems and infection models. *Mathematical Proceedings of the Cambridge Philosophical Society*, 94(01):167–182, 1983.

[FKW13] S. Frischknecht, B. Keller, and R. Wattenhofer. Convergence in (Social) Influence Networks. In *Proc. DISC '13*, pages 433–446, 2013.

[GCFSP85] E. Goles-Chacc, F. Fogelman-Soulié, and D. Pellegrin. Decreasing energy functions as a tool for studying threshold networks. *Discrete Applied Mathematics*, 12(3):261–277, 1985.

[Gif79] D. Gifford. Weighted Voting for Replicated Data. In *Proc. SOSP '79*, pages 150–162, 1979.

[GM90] E. Goles and S. Martínez. *Neural and Automata Networks*. Kluwer, 1990.

[GO80] E. Goles and J. Olivos. Periodic behaviour of generalized threshold functions. *Discrete Mathematics*, 30(2):187–189, 1980.

[GO88] E. Goles and A.M. Odlyzko. Decreasing Energy Functions and Lengths of Transients for Some Cellular Automata. *Complex Systems*, 2(5):501–507, 1988.

[Gol89] E. Goles. Local Graph Transformations Driven by Lyapunov Functionals. *Complex Systems*, 3(1):173–184, 1989.

[HL75] R. Holley and T. Liggett. Ergodic Theorems for Weakly Interacting Infinite Systems and the Voter Model. *The Annals of Probability*, 3(4):643–663, 1975.

[HP01] Y. Hassin and D. Peleg. Distributed Probabilistic Polling and Applications to Proportionate Agreement. *Information and Computation*, 171(2):248–268, 2001.

[KMTN15] D. Kaaser, F. Mallmann-Trenn, and E. Natale. Brief Announcement: On the Voting Time of the Deterministic Majority Process. In *Proc. DISC '15*, 2015.

[KPS13] K. Kothapalli, S. Pemmaraju, and V. Sardeshmukh. On the Analysis of a Label Propagation Algorithm for Community Detection. In *Proc. ICDCN '13*, pages 255–269, 2013.

[KPW14] B. Keller, D. Peleg, and R. Wattenhofer. How Even Tiny Influence Can Have a Big Impact! In *Proc. FUN '14*, pages 252–263, 2014.

[Lig85] T. Liggett. *Interacting Particle Systems*. Springer Berlin Heidelberg, 1985.

[LN07] N. Lanchier and C. Neuhauser. Voter model and biased voter model in heterogeneous environments. *Journal of Applied Probability*, 44(3):770–787, 2007.

[MT14a] F. Mallmann-Trenn. Bounds on the voting time in terms of the conductance. Master's thesis, Simon Fraser University, 2014. Master's thesis. http://summit.sfu.ca/item/14502.

[MT14b] M. Mossel and O. Tamuz. Opinion Exchange Dynamics. *CoRR*, abs/1401.4770, 2014.

[Oli12] R. Oliveira. On the coalescence time of reversible random walks. *Transactions of the American Mathematical Society*, 364(4):2109–2128, 2012.

[Pel02] D. Peleg. Local majorities, coalitions and monopolies in graphs: a review. *Theoretical Computer Science*, 282(2):231–257, 2002.

[Pel14] D. Peleg. Immunity against Local Influence. In *Language, Culture, Computation. Computing - Theory and Technology*, volume 8001 of *LNCS*, pages 168–179. Springer, 2014.

[PS83] S. Poljak and M. Sůra. On periodical behaviour in societies with symmetric influences. *Combinatorica*, 3(1):119–121, 1983.

[PT86] S. Poljak and D. Turzík. On an application of convexity to discrete systems. *Discrete Applied Mathematics*, 13(1):27–32, 1986.

[RAK07] U. Raghavan, R. Albert, and S. Kumara. Near linear time algorithm to detect community structures in large-scale networks. *Physical Review E*, 76(3):036106, 2007.

[SSKc13] A. Sarıyüce, E. Saule, K. Kaya, and U. Çatalyürek. Shattering and Compressing Networks for Betweenness Centrality. In *Proc. SDM '13*, pages 686–694, 2013.

[TT13] O. Tamuz and R. J. Tessler. Majority Dynamics and the Retention of Information. *CoRR*, abs/1307.4035, 2013.

[Win08a] P. Winkler. Puzzled: Delightful Graph Theory. *Communications of the ACM*, 51(8):104, 2008.

[Win08b] P. Winkler. Puzzled: Solutions and Sources. *Communications of the ACM*, 51(9):103, 2008.

Appendix

A. Omitted Proofs from Section 2

Proof of Lemma 2.5. Note that t is at least 4 by definition of the activation time. Let c be the color of K_{black} and $\bar{c} = 1 - c$ the opposite color of c. If at most $m - 2$ of the OR-gates have color $\bar{c}$, then the node of the AND-gate has at least $2+(m-2) > |N(u_0)|/2$ neighbors which will be colored c for all $t \geq 3$ and therefore the AND-gate will be colored c for every $t' \geq 4$.

If, however, $m-1$ of the OR-gates have color $\bar{c}$, only one OR-gate has not been activated and has color c. Thus the node of the AND-gate u_0 has on layer 1 and layer 2 a total of $m - 1$ neighbors of color c and also a total of $m - 1$ neighbors of color $\bar{c}$. That is, these neighbors *cancel* each other out. By Lemma 2.3 the cliques and gates on layers 1 and 2 do not change their color for any $t' \geq 4$. Therefore, the node u_0 can only be influenced by u_1 and the color of u_0 at time t is the color of u_1 at time $t - 1$ for any $t \geq 4$. By Lemma 2.4 we know that u_1 may change its opinion only once in a round $t = a(u_1) \geq 3$ and therefore for any round $t' \geq t + 1$ we have $f_{t'}(u_0) = f_t(u_0)$.

Finally, if m of the OR-gates are colored $\bar{c}$, then u_0 has $m > |N(u_0)|/2$ neighbors of color $\bar{c}$ and since by Lemma 2.3 these m neighbors do not change their color for $t \geq 4$ we have $f_t(u_0) = \bar{c}$ for all $t \geq 4$. Thus, in all cases the claim follows. □

Proof of Lemma 2.6. By Lemma 2.3 all nodes of V_K and V_{OR} are stable after time step 2 and 3, respectively. From Lemma 2.4 we observe that every node of $u_1, \dots, u_k$ with $k = 4 \cdot n$ can only change its color once after time step 3. Note that from Lemma 2.5 we conclude that if u_1 changes its color at time t then the AND-gate does not change its color for any $t' \geq t + 1$.

We now consider the *inner* nodes of the path u_j for which $1 \leq j < k$. In order for a node u_j to change its color at time $t > 3$, one of the neighboring nodes u_{j-1} or u_{j+1} must have changed its color at time $t-1$. This follows, since according to Lemma 2.3 all other neighbors of the node u_j are already stable after 2 steps. Now if a node u_j changes its opinion, one of the neighbors of u_j must have changed its opinion in the previous round. This can only be either u_{j-1} or u_{j+1} (or both), since all other neighbors of u_j are already stable.

Since all nodes $u_1, \dots, u_k$ of the path in G' can only change their color once and since u_0 becomes stable one time step after u_1 changes its color, the convergence time of the graph $G(\Phi)$ is dominated by the behavior of the path. That is, in order to achieve a *long* convergence time, the path must change its color one node after the other, resulting in a convergence time in $\Omega(n)$. Observe that this can only happen if the entire path has a different color than the K_{white} after the second step. As soon as one of the path nodes is assigned the same opinion as the K_{white} mega-clique, the entire path will be activated too early and the process stops prematurely.

Now in order to have a convergence time of $h + 1$, the path in G', $u_0, \dots, u_k$, must activate from u_0 over u_1 up to u_k or in the reverse direction from u_k over u_{k-1} down to

u_0. We now argue that activating from u_k down to u_0 cannot yield a convergence time of at least $h+1$.

Note that all neighbors of u_k except for u_{k-1} are stable at any time step $t \geq 3$. Therefore, u_k either has the same fixed opinion as the K_{white} and thus $a(u_k) = 3$, or u_k has an activation time $a(u_k) = a(u_{k-1}) + 1$. Now in the first case, $a(u_k) = 3$, the convergence time is bounded by $3+k = 3+4n < h+1$, since the path becomes stable one node after the other starting with the node u_k. That is, the resulting convergence time is strictly less than $h+1$. In the second case, $a(u_k) \geq 4$, we note that $a(u_k) = a(u_{k-1}) + 1$ and thus the path cannot activate from u_k down to u_0.

We conclude that in order to have a convergence time of $h+1$ the nodes must activate from u_0 to u_k starting with node u_0 in time step 4 such that $a(u_0) = 4$. Therefore, $a(u_i)$ must be $i+4$ to have a convergence time of $h+1$ which shows the lemma. □

Proof of Lemma 2.7. In the following we use $K^{\mathrm{R}}(x_i)$ and $K^{\mathrm{R}}(\overline{x}_i)$ to denote the representative nodes of the literal cliques for x_i and $\overline{x}_i$. Note that by Lemma 2.3 these representative nodes are stable at time 2. Now assume both cliques have color c after the second step.

Let u be the first node v_1 of the i-th 2/3-gate. By the construction of $G(\Phi)$, u is connected to 6 representative nodes of literal cliques which all share the same color c. Since the representative nodes are stable after 2 steps, u will also have color c for every time step $t' \geq 3$. That is, $a(u) = 3$ and thus by Lemma 2.6 the convergence time is less than $h+1$. □

B. Omitted Proofs from Section 3

Proof of Theorem 3.1. The idea of the proof is to define a potential function ϕ_t that is strictly monotonically decreasing over the time. Let f_0 be any initial opinion assignment. The potential function ϕ_t is simply the number of bad arrows defined in Definition 9, that is

$$\phi_t = \phi_t(G, f_t) = |\{(v,u) \in E : f_{t+1}(u) \neq f_t(v)\}| \quad .$$

Let v denote an arbitrary but fixed node. To show that ϕ_t indeed is a strictly monotonically decreasing potential function as long as $t \leq \mathfrak{T}(G, f_0)$ we distinguish the following two cases.

Case 1. The node v has the same opinion in round $t+1$ as in round $t-1$, that is, $f_{t+1}(v) = f_{t-1}(v)$.

For each neighbor u of v that has a different opinion in round t than v in round $t-1$, there is a bad arrow from v to u. We denote the number of these outgoing bad arrows leaving round $t-1$ as $m_{t-1}(v)$, that is

$$m_{t-1}(v) := |\{u \in N(v) \mid f_t(u) \neq f_{t-1}(v)\}| \quad .$$

There is an incoming bad arrow at node v in round $t+1$ from each neighbor that has a different opinion in round t. Let $n_{t+1}(v)$ be this number, that is

$$n_{t+1}(v) := |\{u \in N(v) \mid f_t(u) \neq f_{t+1}(v)\}| \quad .$$

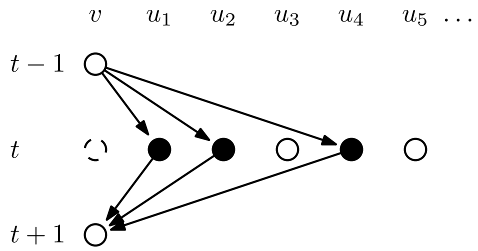

(a) Case 1: $f_{t+1}(v) = f_{t-1}(v)$

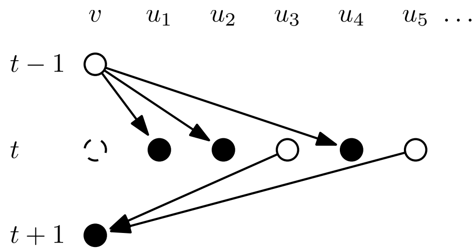

(b) Case 2: $f_{t+1}(v) \neq f_{t-1}(v)$

Figure 4: Above figures show examples for the two cases. In the first case, the number of outgoing bad arrows from v in round $t-1$ equals the number of incoming bad arrows at v in round $t+1$. In the second case the node v has color black in round $t+1$ due to a majority for black in round t. Therefore the number of incoming bad arrows at v in round $t+1$ is strictly smaller than the number of outgoing bad arrows from node v in round $t-1$.

Now recall that v has the same opinion in round $t+1$ as in round $t-1$. Thus, the number of incoming bad arrows at node v in round $t+1$ is the same as the number of bad arrows leaving node v in round $t-1$, which gives us

$$n_{t+1}(v) = m_{t-1}(v) \ . \tag{1}$$

An example for this case is shown in Figure 4a.

Case 2. The node v has a different opinion in round $t+1$ than in round $t-1$, that is, $f_{t+1}(v) \neq f_{t-1}(v)$.

Let $m_{t-1}(v)$ and $n_{t+1}(v)$ be defined as above. Since v changed its opinion after round $t-1$, either in step t or in step $t+1$, there is an incoming bad arrow at node v in round $t+1$ for every neighbor of v that did not have an incoming bad arrow in round t. Now the key is that node v can only have its current opinion in round $t+1$ if there is a clear majority in round t in favor of this opinion among all of its neighbors. Observe that this is where the odd degrees mentioned in the problem statement [Win08a] indeed play a role. Since every node has odd degree, there is always a clear majority among its neighbors and no tie between opinions can ever occur. Now if there is a clear majority in round t, the number of incoming bad arrows at node v in round $t+1$ will be strictly smaller than the number of outgoing bad arrows at node v in round $t-1$, that is

$$n_{t+1}(v) < m_{t-1}(v) \ . \tag{2}$$

An example for this case is shown in Figure 4b.

Both cases. We take the sum over all outgoing bad arrows leaving the nodes in round $t-1$ and obtain $M_{t-1} = \sum_{v \in V} m_{t-1}(v)$. Analogously, we take the sum over all incoming bad arrows in round $t+1$ which gives us $N_{t+1} = \sum_{v \in V} n_{t+1}(v)$. However, since each bad arrow is incident in exactly two nodes, we conclude that the sum over all incoming bad arrows in round $t+1$ is the same as the sum over all outgoing bad arrows in round t. This gives us

$$M_t = N_{t+1} = \phi_t \ . \tag{3}$$

If the deterministic binary majority process has reached a two-periodic state in round t, from Equation 1 and Equation 3 we get

$$\phi_t = N_{t+1} = M_t = \sum_{v \in V} m_t(v) = \sum_{v \in V} n_{t+2}(v) = N_{t+2} = \phi_{t+1} \quad .$$

Now assume that the deterministic binary majority process has not yet reached a two-periodic state in round t. That is, at least one node has a different opinion in round $t+1$ than it had in round $t-1$. Then from Equation 2 and Equation 3 we get

$$\phi_t = N_{t+1} = M_t = \sum_{v \in V} m_t(v) > \sum_{v \in V} n_{t+2}(v) = N_{t+2} = \phi_{t+1}$$

which proves that the voting time of the deterministic binary majority process on G is bounded from above by the initial number of bad arrows.

In particular, since there can be a bad arrow only between ordered pairs of adjacent nodes, the initial number of bad arrows is bounded by $2 \cdot |E|$. Together with the observation that above argument can only be applied after the first step this implies that

$$\mathfrak{T} \leq 2 \cdot |E| + 1 \quad .$$

□

Proof of Theorem 3.2. For every node $v \in V$ the sequence of opinions, $(f_t(v))$, is exactly the same for the deterministic binary majority process in G as in the deterministic binary majority process in G^*. Indeed, every odd-degree node has the same neighborhood in both, G and G^*, thus the process is the same for these nodes. Now consider an arbitrary even-degree node v and fix a round t. If in G there is a tie in round t, v behaves lazily in G and keeps its own opinion at round t. In G^*, the node v considers its own opinion and thus also stays with its own opinion. If on the other hand there is a clear majority in G, this majority has a winning margin of at least 2, since v has even degree. Thus, the impact of the self loop can be neglected and again v behaves the same in G^* as in G.

We can thus bound the voting time of G by applying Theorem 3.1 to the odd-degree graph G^*. □

Proof of Lemma 3.3. From the definition of the deterministic binary majority process we conclude that only less than half of a node's neighbors could have had a different opinion at time $t = 0$, since otherwise the node would have changed its own opinion. Formally, for any $v \in V$ it holds that

$$\sum_{u \in N(v)} [f_1(v) \neq f_0(u)] \leq \frac{|N(v)|}{2} \quad .$$

Also, for odd-degree nodes the above inequality is strict. Therefore, the number of incoming bad arrows at a node at time $t = 1$ is smaller than half of its degree (strictly, for odd nodes). Thus, summing up all initial bad arrows we get $(2 \cdot |E| - |V_{\text{odd}}|)/2$, which concludes the proof. □

B.1. Omitted Proofs from Section 3.1

Definition 12. *In an arbitrary round t an opinion assignment f'_t is a q-*swap *of f_t if for all nodes v*

$$f'_t(v) = f_t(v) \quad \vee \quad f'_t(v) = q \ .$$

That is, all opinions assigned by f'_t are either the original opinion assigned by f_t or q.

Based on this definition we can state and prove the following key lemma.

Lemma B.1 (Monotonicity)**.** *Let f_t be an opinion assignment in round t and f'_t a q-swap of f_t. Let furthermore v be a node for which $f_t(v) \neq f'_t(v)$. It holds for any time step $k \geq t$ that*

$$f_k(v) = q \implies f'_k(v) = q \ .$$

Furthermore, any subsequent opinion assignment f'_k is a q-swap of f_k.

Proof. We show Lemma B.1 by induction over k. The base case for $k = t$ is trivially true. Now suppose that Lemma B.1 holds for $k \leq m$. Let v be an arbitrary but fixed node for which $f_{m+1}(v) = q$. Since $f_{m+1}(v) = q$ we had a majority for q among the neighbors of v in the previous opinion assignment f_m and according to the induction hypothesis f'_m is a q-swap of f_m. Therefore, in f'_m the number of nodes with opinion q could have only increased, strengthening the majority for opinion q even further. Thus, $f'_{m+1}(v) = q$ holds. Now assume f'_{m+1} was not a q-swap of f_{m+1}. That is, there exists a node u for which $f'_{m+1}(u) \neq f_{m+1}(u)$ and $f'_{m+1}(u) \neq q$. This is a contradiction to the previous statement. Together, this concludes the induction. □

In other words, Lemma B.1 states that *strengthening* an opinion will never make it weaker in a subsequent round, that is, if a node ends up with opinion q, it also ends up with the same opinion in the q-swapped opinion assignment.

Definition 13. *An opinion assignment f_t is q-*permanent *if f_{t+2} is a q-swap of f_t.*

We now use the definition above to further bound the voting time, since the deterministic binary majority process has the property that once the process is either in a 0-permanent or 1-permanent state it will converge in a number of steps linear in $|V|$. Furthermore, note that a two-periodic state is both 0-permanent and 1-permanent.

Lemma B.2. *A q-permanent opinion assignment f_t converges in time $2 \cdot |\{v \in V : f_t(v) \neq q\}|$.*

Proof. By definition f_{t+2} is a q-swap of f_t. Thus we can apply Lemma B.1 and conclude that either all nodes have at time $t + 2$ the same opinion as at time t, or some nodes have changed their opinion to q. That is, all nodes with opinion q at time t will also have opinion q at time $t+2$. So there are two possibilities. Either every two time steps at least one node switches to opinion q or every node has again its former opinion and we are in a two-periodic state. Thus the process converges in at most $2\cdot|\{v \in V : f_t(v) \neq q\}| < 2\cdot|V|$ steps. □

We will now use this result to prove an upper bound on the voting time that is better than Corollary 3.4 for dense graphs.

Proof of Theorem 3.5. Let f_t be an opinion assignment that has not yet reached a two-periodic state at time t. In the proof of Lemma B.2 we made the observation that there must exist a node $v \in V$ for which $f_t(v) \neq f_{t-2}(v)$. We therefore distinguish the following two cases.
Case 1. The opinion assignment f_{t-2} is q-permanent.
Case 2. There exists *another* node u with $f_{t-2}(u) \neq f_{t-2}(v)$ such that $f_t(u) \neq f_{t-2}(u)$, that is, u is non-two-periodic and disagrees with v at times t and $t-2$.
As long as we are in case 2, by repeating the argument in case 2 of the proof of Theorem 3.1 we see that the number of bad arrows drops by at least 2 in each step (one due to v and another one due to u). According to Lemma 3.3, this can be the case for at most $1 + \frac{(|E| - |V_{odd}|/2)}{2}$ steps, since the deterministic binary majority process will converge after that time. On the other hand, if at some point we are in case 1, the process will converge in at most $2 \cdot |V|$ steps as shown in Lemma B.2. Together, these two cases yield the bound

$$1 + \frac{|E| - |V_{odd}|/2}{2} + 2 \cdot |V| = 1 + \frac{|E|}{2} - \frac{|V_{odd}|}{4} + \frac{|V|}{4} + \frac{7}{4} \cdot |V| = 1 + \frac{|E|}{2} + \frac{|V_{\text{even}}|}{4} + \frac{7}{4} \cdot |V| \ . \qquad \square$$

B.2. Omitted Proofs from Section 3.2

Proof of Lemma 3.6. Reflexivity and symmetry of $\sim$ hold trivially. What remains to be shown is transitivity, that is, $\forall u, v, w \in V$ it holds that if $\mathrm{fam}(u) = \mathrm{fam}(v)$ and $\mathrm{fam}(v) = \mathrm{fam}(w)$, then also $\mathrm{fam}(u) = \mathrm{fam}(w)$. By definition, we have $N(u) \setminus \{v\} = N(v) \setminus \{u\}$ and $N(v) \setminus \{w\} = N(w) \setminus \{v\}$. By using the previous identities, it follows that $N(u) \setminus \{w, v\} = (N(u) \setminus \{v\}) \setminus \{w\} = (N(v) \setminus \{u\}) \setminus \{w\} = (N(v) \setminus \{w\}) \setminus \{u\} = (N(w) \setminus \{v\}) \setminus \{u\} = N(w) \setminus \{u, v\}$ and

$$v \in N(u) \Leftrightarrow u \in N(v) \Leftrightarrow u \in N(w) \Leftrightarrow w \in N(u) \Leftrightarrow w \in N(v) \Leftrightarrow v \in N(w) \quad . \tag{4}$$

Due to (4), w either belongs to both $N(u)$ and $N(w)$ or to none of them, hence (4) implies $N(u) \setminus \{w\} = N(w) \setminus \{u\}$. This shows transitivity of the relation $\sim$.

From the transitive property of $\sim$ it follows that all nodes in the same family either form a clique or are all non-adjacent. From this latter fact together with the definition

$$\mathrm{fam}(u) = \mathrm{fam}(v) \Leftrightarrow N(u) \setminus \{v\} = N(v) \setminus \{u\}$$

it also follows that all nodes in the same family have the same degree. $\square$

C. Further Computational Properties

In this appendix we investigate some properties of the deterministic binary majority process w.r.t. the potential function of [GO80, PS83], that is, the number of *bad arrows* defined in Definition 9. We show that the convergence time is not monotone w.r.t. the

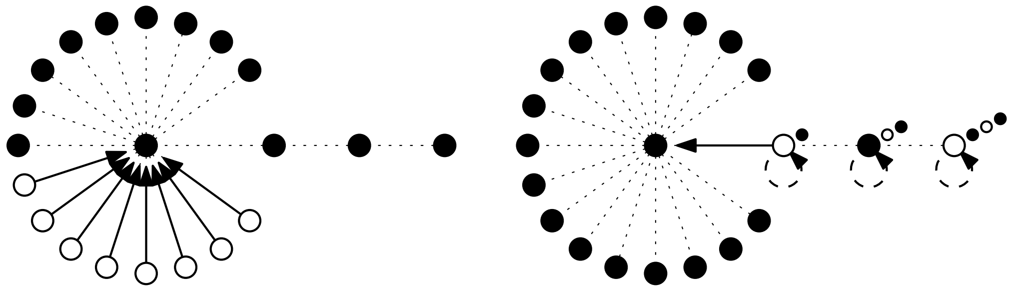

(a) initial opinion assignment $f^{(bad)}$ (b) initial opinion assignment $f^{(good)}$

Figure 5: The figure shows an example for the graph described in the proof of Lemma C.1. It consists of a star graph S_{17} joined at the center node with a path of length 3. Clearly, the initial opinion assignment $f^{(bad)}$ shown in Figure 5a has a total number of 7 bad arrows while the initial opinion assignment $f^{(good)}$ shown in Figure 5a has only one *true* bad arrow along with 3 self loop bad arrows. Still, the process will converge in only one step for $f^{(bad)}$ while it will take 3 steps for $f^{(good)}$.

value of the potential function, and we investigate how many opinion assignments exhibit the same bad arrows. Overall, our results highlight the strengths and weaknesses of such a potential function approach in bounding the voting time of the deterministic binary majority process.

Lemma C.1. *The voting time is not monotone w.r.t. the initial number of bad arrows.*

Proof. Let G be a graph consisting of a star graph S_i with i leaves that has a path graph P_j of length j connected to its center node such that $i > j$. We now can define two initial opinion assignments $f^{(bad)}$ and $f^{(good)}$ for which the initial number of bad arrows in $f^{(bad)}$ is greater than the initial number of bad arrows in $f^{(good)}$ but still $\mathfrak{T}(G, f^{(bad)}) < \mathfrak{T}(G, f^{(good)})$.

As $f^{(bad)}$ assignment, we color $\lceil i/2 \rceil - 1$ leaves of the star graph S_i white and all other nodes, including the path P_j, black. As $f^{(good)}$ assignment, we color all the nodes of S_i black and assign alternating opinions to the nodes of the path P_j. It is straightforward to verify that the described opinion assignments prove the statement. $\square$

An example for a graph G consisting of a S_{17} and a P_3 can be seen in Figure 5. The example shows that even though the initial opinion assignment in Figure 5a has much more initial bad arrows, the deterministic binary majority process converges much faster for the opinion assignment shown in Figure 5b.

Suppose that, instead of specifying the initial opinion assignment, we decide in advance what bad arrows are there. We can do that by deciding for each ordered pair (u, v) for which $\{u, v\} \in E$ whether we want to have a bad arrow going from u to v. We formalize this notion by means of the following definitions.

Definition 14. *Let $G = (V, E)$ be a graph and $\beta : V \times V \to \{0, 1\}$ denote a characteristic function on $V \times V$. Then β is a bad arrows assignment on G if there exists an opinion*

assignment f on G that determines β such that β is the indicator function of the bad arrows we have on G w.r.t. the opinion assignment f.

According to this definition we clearly have $\{u, v\} \notin E \implies \beta(u, v) = 0$ for any bad arrows assignment β. However, there do also exist characteristic functions on the (directed) set of edges of G that do not form a valid bad arrows assignment. An example of such an invalid assignment that motivates above definition is shown in Figure 6.

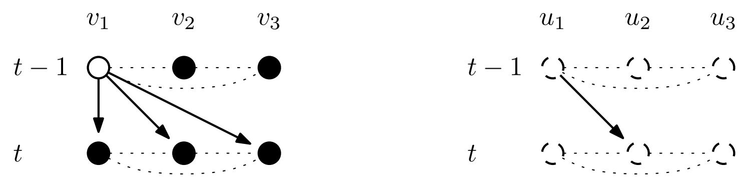

(a) valid assignment (b) invalid assignment

Figure 6: Not every characteristic function on the directed edges is a valid bad arrows assignment.

Figure 6 shows two different assignments of bad arrows for the K_3, a clique of size 3. The left assignment is valid, whereas the right assignment cannot be valid. This is since in cliques of odd size all nodes share the same opinion after exactly one step. Therefore all nodes at step t will have the same opinion. Since, however, u_1 had in step $t-1$ a different opinion than this majority opinion in step t, a bad arrow must exist between u_1 and u_3 (and also a loop from u_1 to itself, if we consider self-loops).

In proving upper bounds on the voting time we consider the bad arrows assignment determined by the initial opinion assignment. One may wonder whether in doing so we are *losing information*. In the following lemma we show that, given a valid bad arrows assignment, we can reconstruct the initial opinion assignment up to exchanging black and white (and up to two more possibilities in bipartite graphs).

Lemma C.2. *Let G be a connected graph and let β be a valid bad arrows assignment on G. If the graph is not bipartite, there are exactly two opinion assignments, otherwise there are exactly four opinion assignments that determine β.*

Proof. Let $v \in V$ denote an arbitrary but fixed vertex. We now denote the set $\{v\}$ as N_0 and the set of direct neighbors of v as N_1 to define the i-th neighborhood N_i for $i \geq 2$ as

$$N_i = \left(\bigcup_{u \in N_{i-1}} N(u) \right) \setminus \left(\bigcup_{j=1}^{i-1} N_j \right) .$$

That is, the set N_i contains all nodes with shortest path to v of length i.

We now show by an induction on $k = 0, 1, 2, \ldots$ that the colors of all nodes in $N_{2 \cdot k}$ are determined by the color of v. The base-case is trivial since for $k = 0$ we have $N_0 = \{v\}$. For the induction step we observe that according to the induction hypothesis the color of each node in $N_{2 \cdot k}$ is determined. We now observe that the color at time 1 of each node in $N_{2 \cdot k+1}$ is determined by β and the colors at time 0 of the nodes in $N_{2 \cdot k}$. Vice versa, also the colors at time 0 of nodes in $N_{2 \cdot (k+1)}$ are determined by β and the colors at time 1 of each node in $N_{2 \cdot k+1}$. This concludes the induction.

An example is shown in Figure 7. In this example it is clear that v and, for example, u_1 must have a different color. Since u_1 does not have a bad arrow to its neighbor in N_1, it has the same color in the next round as this neighbor. But this neighbor's color in the next round is different to the current color of v because of the bad arrow assignment.

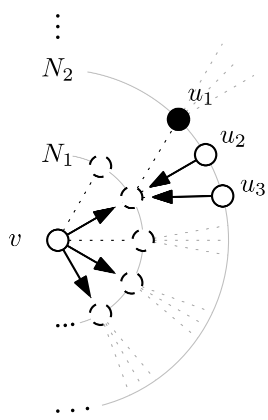

Figure 7: A bad arrows assignment where the opinions of the nodes of each second neighborhood of the graph are uniquely determined.

Observe that from above induction the lemma follows immediately for bipartite graphs. We can fix the colors for two arbitrary nodes, one from each of the two sets of non-adjacent nodes, to determine all other nodes' colors. This gives us four possible opinion assignments for a given bad arrow assignment β. On the other hand, if the graph is not bipartite there must exist a cycle of odd length. The opinion assignments for all nodes of this cycle are determined by β with the same argument as in above induction. Therefore, not only the colors of even neighborhoods $N_{2\cdot k}$ are determined, but also of odd neighborhoods $N_{2\cdot k+1}$. This leaves us with exactly two possible initial opinion assignments, which concludes the proof. $\square$

The following lemma shows that the voting time does not depend, at least straightforwardly, on the diameter.

Lemma C.3. *For any given graph G with diameter Δ, there exists a graph G' with the following properties:*

- *For any opinion assignment f for G, there exists and assignment f' for G' such that the convergence time of G is the same as in G'*
- *The diameter of G' is constant*
- *G is a subgraph of G'.*

Proof. We augment G by adding a clique C_0 of size n where all nodes have Opinion 0 to G. We then add node u_0 initialized with 0 and connect it to all nodes of G and C_0. Symmetrically, we add a clique C_1 of size n where all nodes have Opinion 1 to G. We then add a node u_1 initialized with 1 and connect it to all nodes of G and C_1. Note that every node $u \in G$ is also in G' and the opinion of u is the same in both graphs for any point in time. Hence the convergence time remains the same in G' and the claim follows by observing that G' has a constant diameter. $\square$

Note that above lemma shows that for any connected graph $G = (V, E)$ and any initial opinion assignment f_0 one can construct another graph G' which has G as an induced subgraph, asymptotically the same number of nodes and edges, the same convergence time for a related initial opinion assignment, but a constant diameter. However, there are even examples of graphs where the convergence time of the deterministic binary majority

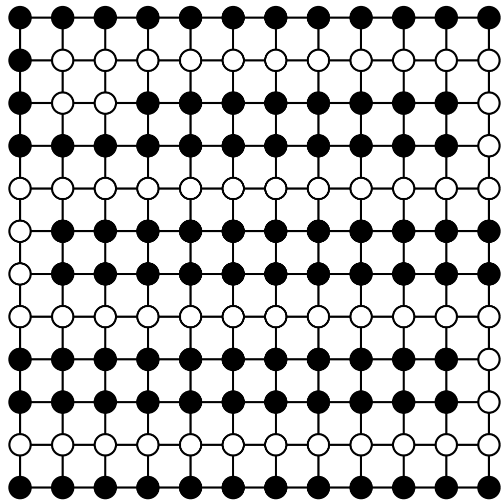

Figure 8: A two-dimensional grid G with $\operatorname{diam}(G) = \sqrt{n}$ and an initial opinion assignment that converges only after $\Omega(n)$ steps.

process w.r.t. a given initial opinion assignment f_0 is asymptotically larger than the diameter of the network without modifying the graph, that is, $\mathfrak{T}(G, f_0) \in \omega(\operatorname{diam}(G))$.

An example for such a graph is shown in Figure 8. In this example, we are given a two-dimensional grid G of size $|V| = \sqrt{n} \times \sqrt{n}$. Clearly, the diameter of this graph is $2 \cdot \sqrt{n}$. However, by laying a winding *serpentine* path of white nodes in an entirely black grid as initial opinion assignment f_0 we can force the process to require a convergence time of $\mathfrak{T}(G, f_0) = \Omega(n) \gg \operatorname{diam}(G) = \mathrm{O}(\sqrt{n})$.

In Theorem 1.2, we show that for the voting time we have $\max_f \mathfrak{T}(G^\Delta, f) \geq \max_f \mathfrak{T}(G, f)$. However, in general it is not the case that $\mathfrak{T}(G^\Delta, f) \geq \mathfrak{T}(G, f)$ for every opinion assignment f, as we show in the following lemma.

Lemma C.4. *Let $G = (V, E)$ be a graph with initial opinion assignment f and G^Δ be the asymmetric graph constructed from G. In general, it does not hold that $\mathfrak{T}(G^\Delta, f) \geq \mathfrak{T}(G, f)$.*

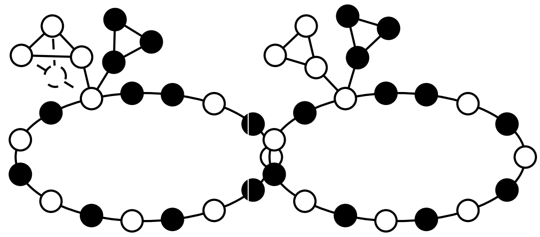

(a) original graph G (b) graph G^Δ

Figure 9: The left graph G is a circle graph with additional gadgets connected to one node. The dashed node is removed to obtain the graph G^Δ shown on the right.

Proof. An example for a graph for which $\mathfrak{T}(G^\Delta, f) \leq \mathfrak{T}(G, f)$ is shown in Figure 9. □